An Efficient Privacy Attack on MyAnimeList's Affinity Oracle: Technical Notes on Score Estimation

Stephen Huan

March 7, 2021

For context (and code) see this GitHub repository.

1 Pearson's Correlation

Definition 1.1. *Pearson's correlation* between vectors $\boldsymbol{u}$ and $\boldsymbol{v}$ is defined as follows: first, define $\boldsymbol{u}' = \boldsymbol{u} - \bar{\boldsymbol{u}}$ where $\bar{\boldsymbol{u}}$ is the constant vector equal to the sample mean of $\boldsymbol{u}$, and define $\boldsymbol{v}'$ similarly. Then Pearson's correlation is simply the cosine similarity:

$$\rho(\boldsymbol{u}, \boldsymbol{v}) = \cos\alpha = \frac{\boldsymbol{u}' \cdot \boldsymbol{v}'}{|\boldsymbol{u}'||\boldsymbol{v}'|}$$

Theorem 1.1. *Pearson's correlation is invariant to a transformation of the form $a\boldsymbol{u}+\boldsymbol{b}$, for any positive real number a and constant vector $\boldsymbol{b}$.*

Proof. We start with the constant vector, which is destroyed when taking the mean:

$$\begin{aligned} \boldsymbol{u}' &= \boldsymbol{u} - \mathrm{E}[\boldsymbol{u}] && \text{definition} \\ (a\boldsymbol{u} + \boldsymbol{b})' &= a\boldsymbol{u} + \boldsymbol{b} - (a\,\mathrm{E}[\boldsymbol{u}] + \boldsymbol{b}) && \text{linearity of expectation} \\ &= a(\boldsymbol{u} - \mathrm{E}[\boldsymbol{u}]) \\ &= a\boldsymbol{u}' \end{aligned}$$

We now see what happens to this scaling by a after Pearson's correlation:

$$\rho(a\boldsymbol{u} + \boldsymbol{b}, \boldsymbol{v}) = \frac{(a\boldsymbol{u}') \cdot \boldsymbol{v}'}{|a\boldsymbol{u}'||\boldsymbol{v}'|} = \frac{a(\boldsymbol{u}' \cdot \boldsymbol{v}')}{a|\boldsymbol{u}'||\boldsymbol{v}'|} = \rho(\boldsymbol{u}, \boldsymbol{v})$$

If $a = 0$ then dividing by the magnitude of the zero vector is ill-defined, and if $a < 0$ then $|a\boldsymbol{u}'|$ becomes $|a||\boldsymbol{u}'|$, so the sign of the correlation switches because of the a on the numerator. Thus, a must be positive for the correlation to be preserved. □

Corollary 1.1.1. *We can't tell the difference between $\boldsymbol{u}$ and $a\boldsymbol{u}+\boldsymbol{b}$ by looking at Pearson's correlation without using additional information.*

2 The Attack

Suppose $\boldsymbol{u}$ is an unknown vector and our only method of obtaining information about $\boldsymbol{u}$ is through querying some *affinity oracle* which tells us $\rho(\boldsymbol{u}, \boldsymbol{v})$ for some $\boldsymbol{v}$ that we pick. We now show how to pick $\boldsymbol{v}$ in such a way that $\boldsymbol{u}$ can be determined, up to a positive scaling and constant translation as mentioned previously.

The basic idea is going to be to pick $\boldsymbol{v}$ such that $\boldsymbol{v}'$ is as close to a basis vector as possible, e.g. something like $\begin{bmatrix}1 & 0 & 0 & \dots & 0\end{bmatrix}^T$. In that case when we compute $\rho(\boldsymbol{u}, \boldsymbol{v})$, it equals $\boldsymbol{u}' \cdot \boldsymbol{v}'$ divided by the magnitude of $\boldsymbol{u}$, which we can ignore since we can't tell the difference between $\boldsymbol{u}$ and $a\boldsymbol{u}$ anyways. Since $\boldsymbol{v}'$ is all zeros except for one 1, $\boldsymbol{u}' \cdot \boldsymbol{v}' = \boldsymbol{u}'_1$, telling us the first element of $\boldsymbol{u}'$ directly. The first snag is that it is not possible for such a $\boldsymbol{v}'$ to exist, since the sum of any $\boldsymbol{x}'$ must be 0 and $\boldsymbol{v}'$ sums to 1.

Theorem 2.1. *The sum of the elements in $\boldsymbol{x}'$ is 0 for any $\boldsymbol{x}$.*

Proof.

$$\mathrm{E}[\boldsymbol{x}'] = \mathrm{E}[\boldsymbol{x} - \mathrm{E}[\boldsymbol{x}]] = \mathrm{E}[\boldsymbol{x}] - \mathrm{E}[\boldsymbol{x}] = 0$$

□

Thus, we need a different pick for $\boldsymbol{v}'$. One simple fix is to introduce a -1 to cancel the existing 1, i.e. $\begin{bmatrix}1 & -1 & 0 & \dots & 0\end{bmatrix}^T$. When we compute $\boldsymbol{u}' \cdot \boldsymbol{v}'$, we get that $\rho(\boldsymbol{u}, \boldsymbol{v}) = \frac{\boldsymbol{u}' \cdot \boldsymbol{v}'}{|\boldsymbol{u}'||\boldsymbol{v}'|} = \boldsymbol{u}_1 - \boldsymbol{u}_2$ (where we arbitrarily let $|\boldsymbol{u}'| = \frac{1}{|\boldsymbol{v}'|}$ for simplicity). This is a linear equation in $\boldsymbol{u}$, so to solve for $\boldsymbol{u}$ we just need a system. The first step is finding a $\boldsymbol{v}$ that generates the $\boldsymbol{v}'$ used above. That can be done by seeing a constant vector will map to the zero vector, so we can start with a vector of all 5's, and then put a 6 where the 1 in $\boldsymbol{v}'$ should be and a 4 where the -1 should be. The average of $\boldsymbol{v}$ will be 5, so $\boldsymbol{v}'$ will map to zeros except the 6 goes to 1 and the 4 goes to -1, which is what we wanted. The second step is that we need a system of equations, so we need multiple (linearly-independent) $\boldsymbol{v}$'s. The obvious choice is to "slide" the 1 over, so our next query is $\begin{bmatrix}0 & 1 & -1 & \dots & 0\end{bmatrix}^T$, the next $\begin{bmatrix}0 & 0 & 1 & -1 & \dots & 0\end{bmatrix}^T$ and so on. Notice if $\boldsymbol{u}$ has M elements, then we only have $M - 1$ versions of $\boldsymbol{v}$ since we can't put a 1 at the last position (we need a -1 after it). This is a natural consequence of the fact that we are missing *two* variables: we can't tell the difference between $\boldsymbol{u}$ and $a\boldsymbol{u} + \boldsymbol{b}$, so a and $\boldsymbol{b}$ are free variables. By setting $|\boldsymbol{u}'|$ we lose a free variable, so we have one free variable left. Thus, our matrix must also have one free variable. We therefore solve the system by putting each variable in terms of the last variable. We have a system of the form:

$$\begin{bmatrix} 1 & -1 & & \\ & 1 & -1 & \\ & & 1 & -1 \end{bmatrix} \boldsymbol{u}' = \boldsymbol{\rho}$$

From the last row, $\boldsymbol{u}'_{n-1} - \boldsymbol{u}'_n = \boldsymbol{\rho}_{n-1}$ so $\boldsymbol{u}'_{n-1} = \boldsymbol{\rho}_{n-1} + \boldsymbol{u}'_n$. We can then add the last row to the second to last row, which cancels the -1 with the 1, yielding $\begin{bmatrix}\dots & 1 & 0 & -1\end{bmatrix}$. Thus, $\boldsymbol{u}'_{n-2} = \boldsymbol{\rho}_{n-2} + \boldsymbol{\rho}_{n-1} + \boldsymbol{u}'_n$. We can then add this row to row above it, and so on.

Thus, the value of $\boldsymbol{u}_i$ is the sum $\boldsymbol{\rho}_i + \boldsymbol{\rho}_{i+1} + \cdots + \boldsymbol{\rho}_{n-1}$ which is a suffix sum that can be computed in $O(n)$ over each element in $\boldsymbol{u}'$. Now that we have $\boldsymbol{u}'$, we need to find a possible $\boldsymbol{u}$ that transforms into it. This could be done directly, see subsection 3.1 for a brute-force analysis, but we'll use a simple observation derived from Theorem 2.1.

Theorem 2.2. $(\boldsymbol{x}')' = \boldsymbol{x}'$, *i.e.* $\boldsymbol{x}'$ *is a fixed point of* $f(\boldsymbol{x}) = \boldsymbol{x} - \mathrm{E}[\boldsymbol{x}]$.

Proof. $(\boldsymbol{x}')' = \boldsymbol{x}' - \mathrm{E}[\boldsymbol{x}']$ by definition. By Theorem 2.1, $\mathrm{E}[\boldsymbol{x}']$ is 0. Thus, $(\boldsymbol{x}')' = \boldsymbol{x}$. □

Recall we wanted to find a $\boldsymbol{x}$ such that $\boldsymbol{x}' = \boldsymbol{u}'$. Since $(\boldsymbol{u}')' = \boldsymbol{u}'$, $\boldsymbol{u}'$ is its own $\boldsymbol{x}$.

2.1 Determining an Valid Integer Form

Now that we have a possible $\boldsymbol{u}$, we need to determine the particular value of a and $\boldsymbol{b}$. One constraint is that each value in a score list must be an integer between 1 and 10, so $a\boldsymbol{u} + \boldsymbol{b} \in \mathbb{Z}^M$ and $1 \leq a\boldsymbol{u} + \boldsymbol{b} \leq 10$. Suppose we had a particular $\boldsymbol{u}$ that satisfies those requirements. Then we could "normalize" $\boldsymbol{u}$ by diving by the greatest common divisor of the score differences, e.g. if $\boldsymbol{u} = \begin{pmatrix}1 & 3 & 1 & 3 & 9 & 9\end{pmatrix}$, we first sort the scores and remove duplicates into $\boldsymbol{u} = \begin{pmatrix}1 & 3 & 9\end{pmatrix}$. We then notice $3-1 = 2$ and $9-3 = 6$, and $\gcd(2,6) = 2$. We can therefore transform $\boldsymbol{u}$ by $\frac{1}{2}\boldsymbol{u}+\frac{1}{2}$ such that 1 goes to 1, 3 goes to 2, and 9 goes to 5. It is now not possible to multiply $\boldsymbol{u}$ by a fraction and then shift $\boldsymbol{u}$ to make all the elements integer because the elements will have different denominators. We sort and only consider the gcd of adjacent elements because if some factor d divides each index i to j, then it must also divide the difference of the element at i from the element at j since the difference is the sum of adjacent differences: $\boldsymbol{v}_j - \boldsymbol{v}_i = (\boldsymbol{v}_{i+1} - \boldsymbol{v}_i) + (\boldsymbol{v}_{i+2} - \boldsymbol{v}_{i+1}) + \cdots + (\boldsymbol{v}_j - \boldsymbol{v}_{j-1})$. Thus, a is now constrained to $\mathbb{Z}$ if $\boldsymbol{u}$ is in this normalized form, and that constrains $\boldsymbol{b} \in \mathbb{Z}^M$ as well. There are now a finite number of possibilities for both, since $a > 0$ and the end result $a\boldsymbol{u} + \boldsymbol{b}$ must be between 1 and 10. Thus, if we have a *single* valid integer representation of $\boldsymbol{u}$, we can normalize it and then generate *all* valid possibilities.

2.2 Finding a Single Valid Integer Form

Recall that we have some $\boldsymbol{u} \in \mathbb{R}^M$ derived from solving a linear system. Since we obtain $\boldsymbol{u}$ by computing the suffix sum of the vector of correlations $\boldsymbol{\rho}$, it will not be integer. We therefore need to find some $\boldsymbol{v}$ such that $\boldsymbol{v} \in \mathbb{Z}^M$, $1 \leq \boldsymbol{v} \leq 10$, and $\boldsymbol{v} = a\boldsymbol{u} + \boldsymbol{b}$ for some positive real number a and constant vector $\boldsymbol{b}$. Framed like this it seems like a search problem, and our search space is exponential: each score has 10 possibilities so a list of length M has 10^M possibilities. In order to cut down on the possibilities, we need to define a *canonical form*. $\boldsymbol{u}$ and $\boldsymbol{v}$ are *equivalent* if there exists a, $\boldsymbol{b}$ such that $v = a\boldsymbol{u} + \boldsymbol{b}$ (the interested reader can verify this defines an equivalence relation, see subsection 3.2). We now need a single, unique, vector to represent the equivalence class defined by all vectors equivalent to some vector $\boldsymbol{u}$. This is the *canonical form* for $\boldsymbol{u}$, or more precisely the canonical form for the equivalence class $\boldsymbol{u}$ belongs to. We can do this simply by "removing" the parameters. We know what the maximum element in $\boldsymbol{u}$ is. If we apply a transformation $a\boldsymbol{u}+\boldsymbol{b}$, then it will still be the maximum element by linearity.

Thus, the maximum and minimum elements give us hints about a. We can divide by $\max(\boldsymbol{u}) - \min(\boldsymbol{u})$ to force the range to be 1, and then shift the resulting vector such that the minimum element is 0 so that $\boldsymbol{b}$ is fixed.

Theorem 2.3. *The canonical form for a vector u is $\mathcal{C}(\boldsymbol{u}) = \frac{\boldsymbol{u} - \min(\boldsymbol{u})}{\max(\boldsymbol{u}) - \min(\boldsymbol{u})}$.*

Proof. We want to show that vectors $\boldsymbol{u}$ and $\boldsymbol{v}$ are equivalent if and only if their canonical forms are equal. We first prove the forward direction, if two vectors are equivalent, then their canonical forms are equal. By hypothesis, $\boldsymbol{u}$ is equivalent to $\boldsymbol{v}$, so $\boldsymbol{v} = a\boldsymbol{u} + \boldsymbol{b}$.

$$\begin{aligned}
\mathcal{C}(\boldsymbol{v}) &= \frac{\boldsymbol{v} - \min(\boldsymbol{v})}{\max(\boldsymbol{v}) - \min(\boldsymbol{v})} \\
&= \frac{a\boldsymbol{u} + b - (a\min(\boldsymbol{u}) + b)}{(a\max(\boldsymbol{u}) + b) - (a\min(\boldsymbol{u}) + b)} \\
&= \frac{a(\boldsymbol{u} - \min(\boldsymbol{u}))}{a(\max(\boldsymbol{u}) - \min(\boldsymbol{u}))} \\
&= \mathcal{C}(\boldsymbol{u})
\end{aligned}$$

We now prove the reverse direction, if $\mathcal{C}(\boldsymbol{u}) = \mathcal{C}(\boldsymbol{v})$ then $\boldsymbol{u}$ and $\boldsymbol{v}$ are equivalent. $\mathcal{C}(\boldsymbol{u})$ is of the form $a\boldsymbol{u} + \boldsymbol{b}$ where $a > 0$, so $\mathcal{C}(\boldsymbol{u}) = \mathcal{C}(\boldsymbol{v})$ implies $a_1\boldsymbol{u} + \boldsymbol{b_1} = a_2\boldsymbol{v} + \boldsymbol{b_2}$. Thus, $\boldsymbol{v} = \frac{a_1}{a_2}\boldsymbol{u} + \frac{1}{a_2}[\boldsymbol{b}_1 - \boldsymbol{b}_2]$, and because $a_2 \neq 0$, $\frac{a_1}{a_2} > 0$, $\boldsymbol{u}$ and $\boldsymbol{v}$ are equivalent. □

We can therefore transform the question "for a given vector $\boldsymbol{u}$, find a vector $\boldsymbol{v}$ such that $\boldsymbol{v} \in \mathbb{Z}^M$, $1 \leq \boldsymbol{v} \leq 10$, and $\boldsymbol{v} = a\boldsymbol{u} + \boldsymbol{b}$" into "for a given vector $\boldsymbol{u}$, find a vector $\boldsymbol{v} \mid \boldsymbol{v} \in \mathbb{Z}^M, 1 \leq \boldsymbol{v} \leq 10$, and $\mathcal{C}(\boldsymbol{u}) = \mathcal{C}(\boldsymbol{v})$". Now we don't have to deal with $\boldsymbol{u}$ at all, just its canonical form, which is helpful since we can construct $\boldsymbol{v}$ from $\mathcal{C}(\boldsymbol{u})$. We first arbitrarily assume $\min(\boldsymbol{v}) = 0$; we can translate $\boldsymbol{v}$ later. Given this assumption, $\mathcal{C}(\boldsymbol{v}) = \frac{\boldsymbol{v}}{\max(\boldsymbol{v})}$. Setting equal to to $\mathcal{C}(\boldsymbol{u})$, $\frac{\boldsymbol{v}}{\max(\boldsymbol{v})} = \mathcal{C}(\boldsymbol{u})$ so $\boldsymbol{v} = \text{round}(\max(\boldsymbol{v})\,\mathcal{C}(\boldsymbol{u}))$ since we need $\boldsymbol{v}$ to be integer. We can then iterate through the possible values for $\max(\boldsymbol{v})$, between 1 and 9 since we assume $\boldsymbol{v}$ is not a constant vector and it is at most 9 greater than the minimum ($10 - 1 = 9$). We choose the maximum value whose computed $\boldsymbol{v}$ is actually equivalent to $\boldsymbol{u}$, measured by the distance between the canonical forms, $|\mathcal{C}(\boldsymbol{u}) - \mathcal{C}(\boldsymbol{v})|$. Ideally this difference should be 0, but because our ρ is noisy (it is only given to 1 decimal place) we need to have some error tolerance. We are guaranteed that at least one $\mathcal{C}(\boldsymbol{v})$ is close to $\mathcal{C}(\boldsymbol{u})$ since there is a ground truth — we got $\boldsymbol{u}$ through Pearson's correlation against a ground truth score vector between 1 and 10.

2.3 Enumerating and Filtering Possibilities

Once we have this $\boldsymbol{v}$, we could still be off from the ground truth by a translation and a scaling. For example, if our vector has scores between 3–10, then we could subtract 2 from each element to yield an equivalent valid integer vector. With just Pearson's, the integer constraint, and the 1–10 constraint we can't distinguish between these two vectors. As mentioned in subsection 2.1, if our vector is sufficiently normalized then a and $\boldsymbol{b}$ are constrained to take integer values (and there is a finite number of pairs). We can

in fact implicitly normalize if in the iteration over $\max(\boldsymbol{v})$ we break ties by prioritizing the smaller value. Given this normalized $\boldsymbol{v}$, we now generate all valid $a, \boldsymbol{b}$ pairs. We simply go through the possible values of a, between 1 and $\lfloor \frac{10}{\max(\boldsymbol{v})} \rfloor$ inclusive, and given a value of a the possible values of $\boldsymbol{b}$ are those which map the smallest to least 1 and the maximum to at most 10, i.e. b can be between $1 - a\min(\boldsymbol{v})$ and $10 - a\max(\boldsymbol{v})$ inclusive.

Finally, we need to guess a particular possibility out of the possible pairs. MAL gives us the private lists's mean to two decimal places, so if the ground truth mean is μ and our mean is $\bar{\boldsymbol{v}}$, then we should pick $a, \boldsymbol{b}$ such that our mean is as close to μ as possible, or when $|a\bar{\boldsymbol{v}} + |\boldsymbol{b}| - \mu|$ is minimized (this follows again from the linearity of expectation).

Another approach is to use maximum likelihood estimation, but that discussion is relegated to the appendix, subsection 3.3 since knowing the sample mean μ is much better than knowing the distribution in practice.

3 Appendix

3.1 Undoing Mean Subtraction

We have a given vector $\boldsymbol{u}'$ and want to find a vector $\boldsymbol{u}$ such that $\boldsymbol{u} - \mathrm{E}[\boldsymbol{u}] = \boldsymbol{u}'$. We can solve for $\boldsymbol{u}$ directly by setting up a linear system, although our linear system will necessarily have one free variable since the transformation is not invertible. The transformation $\boldsymbol{u} - \mathrm{E}[\boldsymbol{u}]$ will be shown to be represented as a matrix A times $\boldsymbol{u}$. This matrix is not invertible because the transformation is not surjective: our matrix is a mapping $\mathbb{R}^M \to \mathbb{R}^M$, and we cannot get vectors whose sum is not 0 by Theorem 2.1. Thus, our function is not surjective and because the domain and co-domain are the same size, our function is not injective: $\boldsymbol{x}$ and $\boldsymbol{x} + \boldsymbol{b}$ for some constant vector $\boldsymbol{b}$ map to the same vector by Theorem 1.1. Thus, our function cannot be invertible on some subspace of $\mathbb{R}^M$. This is all a long-winded way of saying we will have a free variable because we can't distinguish between $\boldsymbol{x}$ and $\boldsymbol{x} + \boldsymbol{b}$. To write $f(\boldsymbol{u}) = \boldsymbol{u} - \mathrm{E}[\boldsymbol{u}]$ as a matrix, we simply write out f.

Taking the first element,

$$\begin{aligned} \boldsymbol{u}_1 - \frac{\boldsymbol{u}_1 + \boldsymbol{u}_2 + \cdots + \boldsymbol{u}_M}{M} &= \boldsymbol{u}_1' \\ \frac{(M-1)\boldsymbol{u}_1 - \boldsymbol{u}_2 - \cdots - \boldsymbol{u}_M}{M} &= \boldsymbol{u}_1' \\ (M-1)\boldsymbol{u}_1 - \boldsymbol{u}_2 - \cdots - \boldsymbol{u}_M &= M\boldsymbol{u}_1' \end{aligned}$$

Thus, we get a matrix of the form:

$$\begin{bmatrix} M-1 & -1 & -1 & -1 \\ -1 & M-1 & -1 & -1 \\ -1 & -1 & M-1 & -1 \\ -1 & -1 & -1 & M-1 \end{bmatrix}$$

If we subtract the second to last row from the last row, we get a row like $\begin{bmatrix} 0 & 0 & \dots & -M & M \end{bmatrix}$ We can then subtract the third to last row from the second to last row, getting a row like $\begin{bmatrix} 0 & \dots & -M & M & 0 \end{bmatrix}$. Repeating the process, we get a matrix of the form:

$$\begin{bmatrix} -M & M & & \\ & -M & M & \\ & & -M & M \end{bmatrix}$$

where we discard the top row since we have a free variable. Amazingly, this matrix is a multiple of the 1 -1 matrix used in section 2, so we have an equation of the form $-MA\boldsymbol{u} = M\mathcal{T}(\boldsymbol{u}')$, so $A\boldsymbol{u} = -\mathcal{T}(\boldsymbol{u}')$. $\mathcal{T}(\boldsymbol{u}')$ is from the fact that we subtracted adjacent rows to force the original matrix into a multiple of A, so $-\mathcal{T}(\boldsymbol{u}')$ looks like $\begin{bmatrix} \boldsymbol{u}_1' - \boldsymbol{u}_2' & \boldsymbol{u}_2' - \boldsymbol{u}_3' & \dots & \boldsymbol{u}_{M-1} - \boldsymbol{u}_M \end{bmatrix}^T$. When we take the suffix sum, we will get $\boldsymbol{u}_i' - \boldsymbol{u}_M'$, so our solution is $\boldsymbol{u}' - \boldsymbol{b}$. This could have been predicted from the start if we realized $\boldsymbol{u}'$ was a constant difference from $\boldsymbol{u}$ (they differ by $\mathrm{E}[\boldsymbol{u}]$) so $\boldsymbol{u}'$ must have been a valid solution to the system to begin with. See archive.py for an implementation, although clearly we just went in a large circle (it's a very elegant circle, however).

3.2 Equivalence Relation Under Pearson's Correlation

Vectors $\boldsymbol{u}$ and $\boldsymbol{v}$ are *equivalent* if they are indistinguishable under Pearson's correlation, i.e. $\rho(\boldsymbol{u}, \boldsymbol{x}) = \rho(\boldsymbol{v}, \boldsymbol{x})$ for all $\boldsymbol{x}$. This is equivalent to saying there exists $a > 0$, $\boldsymbol{b}$ such that $v = a\boldsymbol{u} + \boldsymbol{b}$. Does this define an equivalence relation?

- Reflexivity: Pick $a = 1, b = \mathbf{0}$ so $\boldsymbol{v}$ is equivalent to itself.
- Symmetry. Suppose $\boldsymbol{v} = a\boldsymbol{u} + \boldsymbol{b}$. Then $\boldsymbol{u} = \frac{\boldsymbol{v}-\boldsymbol{b}}{a} = \frac{1}{a}\boldsymbol{v} - \frac{\boldsymbol{b}}{a}$, which is valid for $a \neq 0$.
- Transitivity. Suppose $\boldsymbol{v} = a\boldsymbol{u} + \boldsymbol{b}$, and $\boldsymbol{w} = c\boldsymbol{v} + \boldsymbol{d}$. Then $\boldsymbol{w} = c(a\boldsymbol{u} + \boldsymbol{b}) + \boldsymbol{d} = ca\boldsymbol{u} + [c\boldsymbol{b} + \boldsymbol{d}]$, which is valid since $ca > 0$.

3.3 Maximum Likelihood Estimation

Suppose we have two vectors, $\boldsymbol{u}$ and $\boldsymbol{v}$. Which do we pick? We pick the one that is the most *likely* to occur, assuming we know the distribution. If we have a probability mass function f, $\boldsymbol{u} = \begin{pmatrix}1 & 2 & 2 & 4\end{pmatrix}$, and $\boldsymbol{v} = \begin{pmatrix}1 & 3 & 5 & 5\end{pmatrix}$, then the *likelihood* of $\boldsymbol{u}$ occurring is $\mathcal{L}(\boldsymbol{u}) = f(1) \cdot f(2)^2 \cdot f(4)$ assuming each value of $\boldsymbol{u}$ is identically and independently distributed (i.i.d.). Similarly, the likelihood of $\boldsymbol{v}$ is $\mathcal{L}(\boldsymbol{v}) = f(1) \cdot f(3) \cdot f(5)^2$. We would choose $\boldsymbol{u}$ or $\boldsymbol{v}$ depending on which $\mathcal{L}(\boldsymbol{u})$ or $\mathcal{L}(\boldsymbol{v})$ was bigger. In general,

$$\mathcal{L}(\boldsymbol{x}) = \prod_{i=1}^{M} p(\boldsymbol{x}_i)$$

If we have the frequency of each score x, we can compute this more efficiently as:

$$\mathcal{L}(\boldsymbol{x}) = \prod_{x \in \boldsymbol{x}} p(x)^{\text{count}[x]}$$

Because ln is a monotonic function, we can maximize $\ln \mathcal{L}(\boldsymbol{x})$ instead of $\mathcal{L}(\boldsymbol{x})$, which allows us to turn the product into a sum. This is useful because $\mathcal{L}(\boldsymbol{x})$ decreases exponentially, so we will run into precision issues without the log.

$$\ln \mathcal{L}(\boldsymbol{x}) = \sum_{x \in \boldsymbol{x}} \text{count}[x] \ln p(x)$$

Finally, to pick the a and $\boldsymbol{b}$ which maximize the liklihood, we simply plug $a\boldsymbol{x} + \boldsymbol{b}$ in:

$$\ln \mathcal{L}(a\boldsymbol{x} + \boldsymbol{b}) = \sum_{x \in \boldsymbol{x}} \text{count}[x] \ln p(ax + b)$$